CLAIMS

What is claimed is:

1. A signal bearing medium tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform a method for replacing a failed storage device, the method comprising the following operations:

reconstructing data that was stored on the failed storage device;

commencing writing the reconstructed data on each spare storage device in a plurality of spare storage devices, wherein the reconstructed data is written on the plurality of spare storage devices without verifying that the reconstructed data is successfully written;

finishing writing the reconstructed data on at least one of the spare storage devices in the plurality of spare storage devices;

beginning validating that the reconstructed data was successfully written, on each of the spare storage devices on which writing of the reconstructed data is finished;

receiving a write complete message from at least one of the spare storage devices, indicating that the reconstructed data was successfully written; and

replacing the failed storage device with one of the at least one spare storage devices from which a write complete message was received.

2. The signal bearing medium of claim 1, wherein writing the reconstructed data on the plurality of spare storage devices comprises:

initiating write requests to the plurality of spare storage devices substantially simultaneously; and

synchronizing the heads of the plurality of spare storage devices so the reconstructed data is written substantially simultaneously on the plurality of spare storage devices.

3. The signal bearing medium of claim 1, wherein the operations further comprise lowering a write inhibit threshold for the plurality of spare storage devices before commencing writing the reconstructed data on the plurality of spare storage devices.

4. The signal bearing medium of claim 1, wherein the operations further comprise selecting an alternative seek-settle algorithm for each of the plurality of spare storage devices, wherein each alternative seek-settle algorithm provides faster head movement than the head movement over the same distance when not in a rebuild mode.

5. The signal bearing medium of claim 1, wherein the operations further comprise selecting an alternative settle threshold parameter in a seek-settle algorithm for each of the plurality of spare storage devices, wherein each alternative settle threshold parameter provides a wider settle threshold than the settle threshold for the same head movement distance when not in a rebuild mode.

6. A signal bearing medium tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform a method for replacing a failed storage device in a storage device array, the method comprising the following operations:

detecting that the failed storage device in the storage device array has failed;

reconstructing data that was stored on the failed storage device;

commencing writing the reconstructed data on each spare storage device in a plurality of spare storage devices, wherein the reconstructed data is written on the plurality of spare storage devices without verifying that the reconstructed data is successfully written;

finishing writing the reconstructed data on at least one of the spare storage devices in the plurality of spare storage devices;

beginning validating that the reconstructed data was successfully written, on each of the spare storage devices on which writing of the reconstructed data is finished;

receiving a write complete message from at least one of the spare storage devices, indicating that the reconstructed data was successfully written; and

accepting one of the at least one spare storage devices from which a write complete message was received, into the storage device array to replace the failed storage device.

7. The signal bearing medium of claim 6, wherein the reconstructed data is written on each of the plurality of spare storage devices by initiating write requests to the plurality of spare storage devices substantially simultaneously.

8. The signal bearing medium of claim 7, wherein writing the reconstructed data on the plurality of spare storage devices includes synchronizing the heads of the plurality of spare storage devices.

9. The signal bearing medium of claim 6, wherein the operations further comprise writing new I/O data on at least one spare storage device in the plurality of spare storage devices.

10. The signal bearing medium of claim 6, wherein the spare storage device that is accepted into the storage device array is the first spare storage device from which a write complete message is received.

11. The signal bearing medium of claim 6, wherein the plurality of spare storage devices consists of two spare storage devices.

12. The signal bearing medium of claim 6, wherein the operations further comprise:

designating a plurality of storage devices as the storage device array;

designating a group of spare storage devices for the storage device array, wherein the group of spare storage devices includes at least the plurality of spare storage devices; and

releasing all of the spare storage devices in the plurality of spare storage devices that are not accepted into the storage device array so they may again be used as spare storage devices, after accepting into the storage device array one of the at least one spare storage devices from which a write complete message was received.

13. The signal bearing medium of claim 6, wherein the operations further comprise lowering a write inhibit threshold for the plurality of spare storage devices before commencing writing the reconstructed data on the plurality of spare storage devices.

14. The signal bearing medium of claim 6, wherein the operations further comprise selecting an alternative seek-settle algorithm for each of the plurality of spare storage devices, wherein each alternative seek-settle algorithm provides faster head movement than the head movement over the same distance when not in a rebuild mode.

15. The signal bearing medium of claim 6, wherein the operations further comprise selecting an alternative settle threshold parameter in a seek-settle algorithm for each of the plurality of spare storage devices, wherein each alternative settle threshold parameter provides a wider settle threshold than the settle threshold for the same head movement distance when not in a rebuild mode.

16. A signal bearing medium tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform a method for replacing a failed storage device in a storage device array, the method comprising the following operations:

detecting the failure of the failed storage device in the storage device array;

reconstructing data that was stored on the failed storage device;

commencing writing the reconstructed data on each of a plurality of spare storage devices, wherein the reconstructed data is written on the plurality of spare storage devices without verifying that the reconstructed data is successfully written;

finishing writing the reconstructed data on at least one of the spare storage devices in the plurality of spare storage devices;

beginning validating that the reconstructed data was successfully written, on each of the spare storage devices on which writing of the reconstructed data was finished;

receiving an error message from each of the spare storage devices on which writing of the reconstructed data was finished;

identifying defective data on a first one of the spare storage devices from which an error message was received;

reading replacement data from a second one of the spare storage devices from which an error message was received;

writing the replacement data on the first one of the spare storage devices, to replace the defective data; and

accepting the first one of the spare storage devices into the storage device array to replace the failed storage device.

17. A storage apparatus, comprising:

an array of storage devices;

a plurality of spare storage devices; and

a processor coupled to the array of storage devices and the plurality of spare storage devices, wherein the processor is programmed to perform operations for replacing a failed storage device in the array of data storage devices, the operations comprising:

detecting the failure of the failed storage device in the storage device array;

reconstructing data that was stored on the failed storage device;

commencing writing the reconstructed data on each of the spare storage devices in the plurality of spare storage devices, wherein the reconstructed data is

written on the plurality of spare storage devices without verifying that the reconstructed data is successfully written;

finishing writing the reconstructed data on at least one of the spare storage devices in the plurality of spare storage devices;

beginning validating that the reconstructed data was successfully written, on each of the spare storage devices on which writing of the reconstructed data was finished;

receiving a write complete message from at least one of the spare storage devices, indicating that the reconstructed data was successfully written; and

accepting a first spare storage device from which a write complete message was received, into the storage device array to replace the failed storage device.

18. The storage apparatus of claim 17, wherein the operations further comprise lowering a write inhibit threshold for the plurality of spare storage devices before commencing writing the reconstructed data on the plurality of spare storage devices.

19. The storage apparatus of claim 17, wherein the operations further comprise selecting an alternative seek-settle algorithm for each of the plurality of spare storage devices, wherein each alternative seek-settle algorithm provides faster head movement than the head movement over the same distance when not in a rebuild mode.

20. The storage apparatus of claim 19, wherein the operations further comprise selecting an alternative settle threshold parameter in a seek-settle algorithm for each of the plurality of spare storage devices, wherein each alternative settle threshold parameter provides a wider settle threshold than the settle threshold for the same head movement distance when not in a rebuild mode.

21. A storage system, comprising:

a first memory;

a first device adapter;

a storage device array coupled to the first device adapter;

a plurality of spare storage devices coupled to the first device adapter; and

a first plurality of processors coupled to the first memory and the first device adapter, wherein the first plurality of processors are programmed to perform operations for replacing a failed storage device in a storage device array, the operations comprising:

detecting the failure of the failed storage device in the storage device array;

reconstructing data that was stored on the failed storage device;

writing the reconstructed data simultaneously on each of the plurality of spare storage devices without verifying that that the reconstructed data is successfully written on the plurality of spare storage devices;

finishing writing the reconstructed data on at least one of the spare storage devices in the plurality of spare storage devices;

beginning validating that the reconstructed data was successfully written, on each of the spare storage devices on which writing of the reconstructed data is finished;

receiving a write complete message from at least one of the spare storage devices, indicating that the reconstructed data was successfully written; and

accepting a first spare storage device from which a write complete message was received, into the storage device array to replace the failed storage device.

22. The storage system of claim 21, further comprising:

a second memory;

a second device adapter coupled to the storage device array and the plurality of spare storage devices;

a second plurality of processors coupled to the second memory and the second device adapter; and

a plurality of shared adapters coupled to the first plurality of processors and the second plurality of processors.

23. The storage system of claim 22, wherein the operations further comprise:

lowering a write inhibit threshold for the plurality of spare storage devices before commencing writing the reconstructed data on the plurality of spare storage devices;

selecting an alternative seek-settle algorithm for each of the plurality of spare storage devices, wherein each alternative seek-settle algorithm provides faster head movement than the head movement over the same distance when not in a rebuild mode; and

selecting an alternative settle threshold parameter in a seek-settle algorithm for each of the plurality of spare storage devices, wherein each alternative settle threshold parameter provides a wider settle threshold than the settle threshold for the same head movement distance when not in a rebuild mode.

24. A system for reducing the rebuild time of a storage device array, comprising:

means for detecting a failed storage device in the storage device array;

means for reconstructing data that was stored on the failed storage device;

means for commencing writing the reconstructed data on each of a plurality of spare storage devices, wherein the reconstructed data is written on the plurality of spare storage devices without verifying that the reconstructed data is successfully written;

means for finishing writing the reconstructed data on at least one of the spare storage devices in the plurality of spare storage devices;

means for beginning validating that the reconstructed data was successfully written, on each of the spare storage devices on which writing of the reconstructed data is finished;

means for receiving a write complete message from at least one of the spare storage devices, indicating that the reconstructed data was successfully written; and

means for accepting one of the at least one spare storage devices from which a write complete message was received, into the storage device array to replace the failed storage device.

25. A method for replacing a failed storage device in a storage device array, comprising:

detecting the failure of the failed storage device in the storage device array;

reconstructing data that was stored on the failed storage device;

commencing writing the reconstructed data on each of a plurality of spare storage devices, wherein the reconstructed data is written on the plurality of spare storage devices without verifying that the reconstructed data is successfully written;

finishing writing the reconstructed data on at least one of the spare storage devices in the plurality of spare storage devices;

beginning validating that the reconstructed data was successfully written, on each of the spare storage devices on which writing of the reconstructed data is finished;

receiving a write complete message from at least one of the spare storage devices, indicating that the reconstructed data was successfully written; and

accepting one of the at least one spare storage devices from which a write complete message was received, into the storage device array to replace the failed storage device.

26. The method of claim 25, wherein writing the reconstructed data on the plurality of spare storage devices comprises:

initiating write requests to the plurality of spare storage devices substantially simultaneously; and

synchronizing the heads of the plurality of spare storage devices so the reconstructed data is written substantially simultaneously on the plurality of spare storage devices.

27. The method of claim 25,

further comprising writing new I/O data on the plurality of spare storage devices; and

wherein the spare storage device accepted into the storage device array is the first spare storage device from which a write complete message is received.

28. The method of claim 26, further comprising:

lowering a write inhibit threshold for the plurality of spare storage devices before commencing writing the reconstructed data on the plurality of spare storage devices;

selecting an alternative seek-settle algorithm for each of the plurality of spare storage devices, wherein each alternative seek-settle algorithm provides faster head movement than the head movement over the same distance when not in a rebuild mode; and

selecting an alternative settle threshold parameter in a seek-settle algorithm for each of the plurality of spare storage devices, wherein each alternative settle threshold parameter provides a wider settle threshold than the settle threshold for the same head movement distance when not in a rebuild mode.